



The
Patent
Office

PC/GB 97/00060

09/043171

The Patent Office
Cardiff Road
Newport
Gwent
NP9 1RH

PRIORITY DOCUMENT

REC'D 14 FEB 1997

WIPO

PCT

I, the undersigned, being an officer duly authorised in accordance with Section 74(1) and (4) of the Deregulation and Contracting Out Act 1994, to sign and issue certificates on behalf of the Comptroller-General, hereby certify that annexed hereto is a true copy of the documents as originally filed in connection with the patent application identified therein.

In accordance with the Patents (Companies Re-registration) Rules 1982, if a company named in this certificate and any accompanying documents has re-registered under the Companies Act 1980 with the same name as that with which it was registered immediately before re-registration save for the substitution as, or inclusion as, the last part of the name of the words "public limited company" or their equivalents in Welsh, references to the name of the company in this certificate and any accompanying documents shall be treated as references to the name with which it is so re-registered.


In accordance with the rules, the words "public limited company" may be replaced by p.l.c., plc, P.L.C. or PLC.

Re-registration under the Companies Act does not constitute a new legal entity but merely subjects the company to certain additional company law rules.

Signed

Dated

7 FEB 1997



5 JAN 1996

Request for grant of a patent

(See the notes on the back of this form. You can also get an explanatory leaflet from the Patent Office to help you fill in this form)

The Patent Office

Cardiff Road
Newport
Gwent NP9 1RH

1. Your reference

J23735 GB

2. Patent application number

(The Patent Office will fill in this part)

9600774.5

3. Full name, address and postcode of the or of each applicant (underline all surnames)

British Telecommunications Public Limited Company
81 Newgate Street
London, EC1A 7AJ
United Kingdom

Patents ADP number (if you know it)

If the applicant is a corporate body, give the country/state of its incorporation

18767002es

4. Title of the invention

WAVEFORM SYNTHESIS

5. Name of your agent (if you have one)

"Address for service" in the United Kingdom to which all correspondence should be sent (including the postcode)

R.G.C. Jenkins & Co.
26 Caxton Street
London, SW1H 0RJ
United Kingdom

Patents ADP number (if you know it)

950001

6. If you are declaring priority from one or more earlier patent applications, give the country and the date of filing of the or of each of these earlier applications and (if you know it) the or each application number

Country

Priority application number
(if you know it)

Date of filing
(day / month / year)

7. If this application is divided or otherwise derived from an earlier UK application, give the number and the filing date of the earlier application

Number of earlier application

Date of filing
(day / month / year)

8. Is a statement of inventorship and of right to grant of a patent required in support of this request? (Answer 'Yes' if:

YES

- a) any applicant named in part 3 is not an inventor, or
 - b) there is an inventor who is not named as an applicant, or
 - c) any named applicant is a corporate body.
- See note (d))

Patents Form 1/77

9. Enter the number of sheets for any of the following items you are filing with this form. Do not count copies of the same document

Continuation sheets of this form

Description	22
Claim(s)	3
Abstract	1
Drawing(s)	12

10. If you are also filing any of the following, state how many against each item.

Priority documents

Translations of priority documents

Statement of inventorship and right to grant of a patent (Patents Form 7/77)

Request for preliminary examination and search (Patents Form 9/77)


Request for substantive examination (Patents Form 10/77)

Any other documents (please specify)

11.

I/We request the grant of a patent on the basis of this application.

Signature



Date 15.01.96

R.G.C. JENKINS & CO.

- 12 Name and daytime telephone number of person to contact in the United Kingdom David C. Musker 0171- 931 7141

Warning

After an application for a patent has been filed, the Comptroller of the Patent Office will consider whether publication or communication of the invention should be prohibited or restricted under Section 22 of the Patents Act 1977. You will be informed if it is necessary to prohibit or restrict your invention in this way. Furthermore, if you live in the United Kingdom, Section 23 of the Patents Act 1977 stops you from applying for a patent abroad without first getting written permission from the Patent Office unless an application has been filed at least 6 weeks beforehand in the United Kingdom for a patent for the same invention and either no direction prohibiting publication or communication has been given, or any such direction has been revoked.

Notes

- If you need help to fill in this form or you have any questions, please contact the Patent Office on 0645 500505.
- Write your answers in capital letters using black ink or you may type them.
- If there is not enough space for all the relevant details on any part of this form, please continue on a separate sheet of paper and write "see continuation sheet" in the relevant part(s). Any continuation sheet should be attached to this form.
- If you have answered 'Yes' Patents Form 7/77 will need to be filed.
- Once you have filled in the form you must remember to sign and date it.
- For details of the fee and ways to pay please contact the Patent Office.

WAVEFORM SYNTHESIS

This invention relates to methods and apparatus for waveform synthesis, and particularly but not exclusively for speech synthesis.

5 Various types of speech synthesizer are known. Most operate using a repertoire of phonemes or allophones, which are generated in sequence to synthesise corresponding utterances. A review of some types of speech synthesizers may be found in A. Breen "Speech Synthesis Models: A Review",
10 Electronics and Communication Engineering Journal, pages 19-31, February 1992. Some types of speech synthesizer attempt to model the production of speech by using a source-filter approximation utilising, for example, linear prediction. Others record stored segments of actual speech, which are
15 output in sequence.

A major difficulty with synthesised speech is to make the speech sound natural. There are many reasons why synthesised speech may sound unnatural. However, a particular problem with the latter class of speech synthesizers, utilising
20 recorded actual speech, is that the same recording of each vowel or allophone is used on each occasion where the vowel or allophone in question is required. This becomes even more noticeable in those synthesizers where, to generate a sustained sound, a short segment of the phoneme or allophone
25 is repeated several times in sequence.

The present invention, in one aspect, provides a speech synthesizer in which a speech waveform is directly synthesised by selecting a synthetic starting value and then selecting and outputting a sequence of further values, the selection of each
30 further value being based jointly upon the value which preceded it and upon a model of the dynamics of actual recorded human speech.

Thus, a synthesised sequence of any required duration can be generated. Furthermore, since the progression of the
35 sequence depends upon its starting value, different sequences corresponding to the same phoneme or allophone can be generated by selecting different starting values.

The present inventors have previously reported ("Speech characterisation by non-linear methods", M. Banbrook and S. McLaughlin, submitted to IEEE Transactions on Speech and Audio Processing, 1996; "Speech characterisation by non-linear methods", M. Banbrook and S. McLaughlin, presented at IEEE Workshop on non-linear signal and image processing, pages 396-400, 1995) that voiced speech, with which the present invention is primarily concerned, appears to behave as a low dimensional, non-linear, non-chaotic system. Voiced speech is essentially cyclical, comprising a time series of pitch pulses of similar, but not identical, shape. Therefore, in a preferred embodiment, the present invention utilises a low dimensional state space representation of the speech signal, in which successive pitch pulse cycles are superposed, to estimate the progression of the speech signal within each cycle and from cycle-to-cycle.

This estimate of the dynamics of the speech signal is useful in enabling the synthesis of a waveform which does not correspond to the recorded speech on which the analysis of the dynamics were based, but which consists of cycles of a similar shape and exhibiting a similar variability to those on which the analysis was based.

For example, the state space representation may be based on Takens' Method of Delays (F. Takens, "Dynamical Systems and Turbulence", Vol. 898 of Lecture Notes in Mathematics, pages 366-381. Berlin: Springer 1981). In this method, the different axes of the state space consist of waveform values separated by predetermined time intervals, so that a point in state space is defined by a set of values at t_1 , t_2 , t_3 (where $t_2 - t_1 = \Delta_1$ and $t_3 - t_2 = \Delta_2$, which are both constants and may be equal).

Another current problem with synthesised speech is that where different sounds are concatenated together into a sequence, the "join" is sometimes audible, giving rise to audible artifacts such as a faint modulation at the phoneme rate in the synthesised speech.

Accordingly, in another aspect the present invention

provides a method and apparatus for synthesising speech in which an interpolation is performed between state space representations of the two speech sounds to be concatenated, or, in general, between correspondingly aligned portions of each pitch period of the two sounds. Thus, one pitch pulse shape is gradually transformed into another.

Other aspects and preferred embodiments of the invention will be apparent from the following description and claims.

The invention will now be illustrated, by way of example only, with reference to the accompanying drawings in which:-

Figure 1 is a diagram of signal amplitude against time for a (notional) voiced speech signal;

Figure 2 is a diagram of signal amplitude against time for a notional cyclical waveform, illustrating the derivation of state sequence points based on the method of delays;

Figure 3 is a state sequence space plot of the points of Figure 2;

Figure 4 is a state sequence space plot showing the trajectory of a notional voiced speech sound defining an attractor in the state sequence space;

Figure 5 is an illustrative diagram, on a formant chart showing state sequence space attractors (corresponding to that of Figure 4) for a plurality of different vowels;

Figure 6 is a block diagram showing schematically the structure of a speech synthesizer according to a first embodiment of the invention;

Figure 7 is a flow diagram showing illustratively the method of operation of the speech synthesizer of Figure 6;

Figure 8 is a time line showing illustratively the sequence of speech and silence segments making up a speech utterance;

Figure 9a is a state sequence space plot showing a single cycle of a notional voiced sound, and a portion of a cycle of a synthesised sound synthesised therefrom;

Figure 9b is a detail of Figure 9a;

Figure 9c is a state sequence space diagram showing multiple cycles of a waveform, and

Figure 9d is a detail thereof showing the neighbourhood surrounding a point on one cycle, the transformation of which over time is utilised in the embodiment of Figure 6;

Figure 10 is a block diagram showing schematically the components of apparatus for deriving the synthesised data used in the embodiment of Figure 6;

Figures 11a-d illustrates the data produced at various stages of the process of operation of the apparatus of Figure 10;

Figure 12 is a flow diagram illustrating the stages of operation of the apparatus of Figure 10;

Figure 13 is a state sequence space diagram showing illustratively the effect of the transformation over time of the neighbourhood of Figure 9c;

Figure 14 is a flow diagram showing in greater detail the process of progressing from one sound to another forming part of the flow diagram of Figure 7;

Figure 15 is an illustrative diagram indicating the combination of two state space sequences performed during the process of Figure 14; and

Figure 16 is a flow diagram showing the process of progressing from one sound to another in a second embodiment of the invention.

State space representation of the speech signal

Before describing embodiments of the invention in detail, a brief description will be given of the state space representation of the speech signal utilised in embodiments of the invention (but known in itself as a tool for analysis of speech, for example from "Lyapunov exponents from a time series: a noise-robust extraction algorithm"; M. Banbrook, G. Ushaw, S. McLaughlin, submitted to IEEE Transactions on signal processing, October 1995 to which reference is made if further detail is required).

Figure 1 illustrates a speech signal or, more accurately, a portion of a voiced sound comprised within a speech signal. The signal of Figure 1 may be seen to consist of a sequence of similar, but not identical, pitch pulses p_1 , p_2 , p_3 . The

shape of the pitch pulses characterises the sound of the phoneme, and their period characterises the pitch perceived.

Referring to Figure 2, to produce a state space representation of a time sequence X , a plurality (in this case 5 3) of values of the waveform X at spaced apart times x_{i-10} , x_i , x_{i+10} are taken and combined to represent a single point \underline{s}_i in a space defined by a corresponding number of axes.

Thus, referring to Figures 2 and 3, a first point \underline{s}_1 is represented by the three dots on the curve X representing 10 values of the waveform X at sample times x_0 , x_{10} and x_{20} . Since all three of these values are positive, the point they define \underline{s}_1 lies in the positive quadrant of the space of Figure 3.

A further point \underline{s}_2 is represented by the three crosses 15 in Figure 2 on the waveform X . This point is defined by the three values x_1 , x_{11} and x_{21} . Since all three of these values are more positive than those of the point \underline{s}_1 , the point \underline{s}_2 in the state sequence space of Figure 3 will lie in the same quadrant and radially further out than the point \underline{s}_1 .

20 Likewise, a third point \underline{s}_3 is defined by values of the waveform X at times 2, 12 and 22 (x_2 , x_{12} and x_{22} respectively). This point is indicated by three triangles on the waveform X in Figure 2.

Thus, in general, in this time delay method of 25 constructing a state space representation of the time sequence X (i.e. speech waveform), for each successive time sample X_i , the corresponding point \underline{s}_i in the state sequence space is represented by the value of that point X_i together with those of a preceding and a succeeding point x_{i-j} , x_{i+k} (where j is 30 conveniently equal to k and in this case both are equal to 10).

If the waveform of Figure 2 were simply a diagonal straight line, its representation in the state space of Figure 3 would likewise be a straight line.

35 However, for a repetitive time sequence of the type shown in Figures 1 or 2, points of inflection in the waveform cause the corresponding sequence of points in state space to define

a trajectory which likewise inflects, and follows a closed loop to return to close to its start point. Since the relative values of the points x_i , x_{i-j} , x_{i+k} , depend little on time and are closely similar for successive cycles of the time sequence they represent, referring to Figure 4, the state space representation of a sequence of N cycles (e.g. pitch pulses p_1-p_n) of a waveform will be a continuous trajectory through the state sequence space performing N closely similar circuits, so as to define a circuitous multidimensional surface, or manifold, containing N strands or tracks. The surface which would be generated by an infinite number of such cycles is referred to as the "attractor" of the waveform X giving rise to it.

The attractor of Figure 4 consists of a double loop (which, in the projection indicated, appears to cross itself but does not in fact do so in three dimensions).

Referring to Figure 5, we have determined that each voiced sound gives rise to an attractor of this nature, all of which can adequately be represented in a three dimensional state space, although it might also be possible to use as few as two dimensions or as many as four, five or more. The important parameters for an effective representation of voiced sounds in such a state space are the number of dimensions selected and the time delay between adjacent samples.

As shown in Figure 5 in which the axes over which the attractors are distributed are f_1 (the frequency of the first formant) against f_2-f_1 (where f_2 is the frequency of the second formant), the shapes of the attractors vary considerably (with the corresponding shapes of the speech waveforms to which they correspond) although there is some relationship between the topologies of respective attractors and the sounds to which they correspond.

The discussion above relates to voiced sounds (such as vowels and voiced consonants). It is, of course, possible to provide a state sequence representation of any waveform, but in the case of unvoiced sounds (i.e. fricatives) the state space representation will not follow successive closely

similar loops with a well defined topology, but instead will follow a trajectory which passes in an apparently random fashion through a volume in the state sequence space.

Overview of first embodiment of the invention

5 Referring to Figure 6, in a first embodiment of the invention a speech synthesizer comprises a loudspeaker 2, fed from the analogue output of a digital to analog converter 4, coupled to an output port of a central processing unit 6 in communication with a storage system 8 (comprising random
10 access memory 8a, for use by the CPU 6 in calculation; program memory 8b for storing the CPU operating program; and data constant memory 8c for storing data for use in synthesis).

The apparatus of Figure 6 may conveniently be provided by a personal computer and sound card such as an Elonex (TM)
15 Personal Computer comprising a 33 mHz Intel 486 microprocessor as the CPU 6 and an Ultrasound Max (TM) soundcard providing the digital to analogue converter 4 and output to a loudspeaker 2. Any other digital processor of similar or higher power could be used instead.

20 Conveniently, the storage system 8 comprises a mass storage device (e.g. a hard disk) containing the operating program and data to be used in synthesis and a random access memory comprising partitioned areas 8a, 8b, 8c, the program and data being loaded into the latter two of which,
25 respectively, prior to use of the apparatus of Figure 6.

The stored data held within the stored data memory 8c comprises a set of records 10a, 10b, ... 10c, each of which represents a small segment of a word which may be considered to be unambiguously distinguishable regardless of its context
30 in a word or phrase (i.e. each corresponds to a phoneme or allophone). The phonemes can be represented by any of a number of different phonetic alphabets; in this embodiment, the SAMPA (Speech Assessment Methodology Phonetic Alphabet, as disclosed in A. Breen, "Speech Synthesis Models: A Review",
35 Electronics and Communication Engineering Journal, pages 19-31, February 1992) is used. Each of the records comprises a respective waveform recording 11 comprising successive digital

values (e.g. sampled at 20 kHz) of the waveform of an actual utterance of the phoneme in question as successive samples x_1 , $x_2 \dots x_N$.

Additionally, each of the records 10 associated with a
5 voiced sound (i.e. the vowels and voiced consonant sounds of the phonetic alphabet) comprises, for each stored sample x_i , a transform matrix defined by nine stored constant values.

Thus, the data memory 8c comprises on the order of thirty
10 to forty records 10 (depending the phonetic alphabet chosen), each consisting of on the order of half a second of recorded digital waveforms (i.e., for sampling at 20 kHz, around ten thousand samples x_i , each of the sample records for voiced sounds having an associated nine element transform matrix). The volume required by the data memory 8c is thus ((9+1) x
15 10,000 x 40 = 400,000) 16 bit memory locations.

The manner in which the contents of the data memory 8c are derived will be described in greater detail below.

As indicated in Figure 8, an utterance to be synthesised by the speech synthesizer consists of a sequence of portions
20 each with an associated duration, comprising a silence portion 14a followed by a word comprising a sequence of portions 14b-14f each consisting of a phoneme of predetermined duration, followed by a further silence portion 14g, followed by a further word comprised of phoneme portions 14h-14j each of an
25 associated duration, and so on. The sequence of phonemes, together with their durations, are either stored or derived by one of several well known rule systems forming no part of the present invention, but comprised within the control program.

30 Referring to Figure 7, the operation of the control program of the CPU 6 will now be described in greater detail.

In accordance with a sequence thus determined, in a step 502, the CPU 6 selects a first sound record 10 corresponding to one of the phonemes of the sequence illustrated in Figure
35 8.

In a step 504, the CPU 6 executes a transition to the sound as will be described in greater detail below.

In a step 506, the CPU 6 selects a start point for synthesis of the phoneme waveform, \hat{x}_i . Referring to Figure 9, the selection of the start point for synthesis consists of two stages. Firstly, as a result of the progression step 504, as discussed in greater detail below, the CPU 6 will have selected some point x_i on the stored waveform. The next step is then to select a new point, randomly located in the state sequence space, but close to the already selected point.

For example, referring to Figure 9b, the most recent stored point accessed by the CPU 6 (and output to the DAC 4 and hence the loudspeaker 2 as synthesised sound) is point x_{21} with corresponding state space point \underline{s}_{21} , and in step 506, a first synthesised start point $\hat{\underline{s}}_1$ is selected close to x_{21} .

The mechanism for selecting a close point may be as follows:

1. The first point \underline{s}_i in state sequence space is found by reading values x_i , x_{i-10} and x_{i+10} .
2. The next point \underline{s}_{i+1} on the trajectory in state sequence space is found by accessing values x_{i+1} , x_{i+11} , and x_{i-9} .
3. The euclidean (i.e. root mean square) distance in the state sequence space between the two points \underline{s}_i , \underline{s}_{i+1} is calculated.
4. A pseudo random sequence algorithm is used to generate the random coordinates of a point $\hat{\underline{s}}_1$ in state space, spaced from the point \underline{s}_i by a euclidean distance between zero and the distance thus calculated.

Having thus determined a first synthesised start point $\hat{\underline{s}}_1$ close to, but not coincident with, the stored trajectory, in the region of the last actual point output (x_{21} in this case), in step 508, the CPU 6 determines the closest point on the stored trajectory to the newly synthesised point $\hat{\underline{s}}_1$.

Very often the closest point selected in step 508 will in fact be the last point on the stored trajectory (in this case \underline{s}_{21}). However, it may correspond instead to one of the nearest neighbours on the stored trajectory (as in this case, where \underline{s}_{22} is closer), or to a point on another strand of the stored trajectory where this is closely spaced in the state

sequence space, as indicated in Figure 9c.

Having thus determined the closest point on the stored trajectory made up of the stored waveform points x_i , the CPU 6 is arranged in step 510 to calculate the offset vector from the closest point on the stored trajectory thus selected in step 508 to the synthesised point \hat{s}_i . The offset vector b_i thus calculated therefore comprises a three element vector.

Next, in step 512, the next offset vector b_{i+1} (in this case b_2) is calculated by the CPU 6, by reading the matrix T_i stored in relation to the preceding point x_i (in this case in relation to point x_{22}) and multiplying this by the transpose of the first offset vector b_i (in this case b_1).

Next, in step 514, the CPU 6 selects the next stored trajectory point s_{i+1} , in this case, point s_{23} (defined by values x_{23} , x_{13} and x_{33}).

In step 516, the next synthesised speech point is calculated (\hat{s}_{i+1}) by adding the newly calculated offset vector b_{i+1} to the next point on the trajectory s_{i+1} . Then, the centre value x_{i+1} of the newly synthesised point \hat{s}_{i+1} is output to the DAC 4 and loudspeaker 2.

In step 520, the CPU 6 determines whether the required predetermined duration of the phoneme being synthesised has been reached. If not, then the CPU 6 returns to step 508 of the control program, and determines the new closest point on the trajectory to the most recently synthesized point. In many cases, this may be the same as the point s_{i+1} from which the synthesised point was itself calculated, but this is not necessarily so.

Thus, by following the process of steps 506-518, the CPU 6 is able to synthesis a speechlike waveform (shown as a dashed trajectory in state sequence space in Figures 9a and 9b) from the stored waveform values x_i and transform matrices T_i .

The length of the synthesised sequence does not in any way depend upon the number of stored values, nor does the synthesised sequence exactly replicate any portion of the stored sequence.

Instead, each point on the synthesised sequence depends jointly upon the preceding point in the synthesised sequence; the nearest other points (in state sequence space) in the stored sequence; and the transform matrix in relation to the
 5 nearest point in the stored sequence.

Thus, due to the random selection of start points in step 506, the synthetic waveform generated will differ from one synthesis process to the next.

When the predetermined end point for the phoneme in
 10 question has been reached in step 520, in step 522 the CPU 6 determines whether the end of the desired sequence (e.g. as shown in Figure 8) has been reached, and if so, in a step 524 the CPU 6 causes the output sequence to progress to silence (as will be discussed in greater detail below).

15 If not, the CPU 6 selects the next sound in the sequence (step 525) and determines, in a step 526, whether the next sound is voiced or not. If the next sound is voiced, the CPU 6 returns to step 502 of Figure 7, whereas if the next sound is unvoiced, in a step 528 the CPU 6 progresses (as will be
 20 described in greater detail below) to the selected unvoiced sound, which is then reproduced in step 530 (as will be described in greater detail below). The CPU 6 then returns to step 522 of Figure 7.

Calculation of transform matrix

25 Referring to Figure 10, apparatus for deriving the stored sample and transform records 10 comprises a microphone 22, an analog to digital converter 24, a CPU 26, and a storage device 28 (provided, for example, by a mass storage device such as a disk drive and random access memory) comprising a working
 30 scratch pad memory 28a and a program memory 28b.

Naturally, the CPU 26 and storage device 28 could be physically comprised by those of a speech synthesizer as shown in Figure 6, but it will be apparent that this need not be the case since the data characterising the speech synthesizer of
 35 Figure 6 is derived prior to, and independently of, the synthesis process.

Conveniently, the analog to digital converter 24 is

arranged to sample the analog speech waveform from the microphone 22 at a frequency of around 20 kHz and to an accuracy of 16 bits.

Referring to Figures 11 and 12, the operation of the apparatus of Figure 10 will now be described. In a step 602, as shown in Figure 11a, whilst a human speaker recites a single utterance of a desired sound (e.g. a vowel) the CPU 26 and analog to digital converter 24 sample the analog waveform thus produced at the output of the microphone 22 and store successive samples (e.g. around 10,000 samples, corresponding to around half a second of speech) in the working memory area 28a.

Next, in a step 604, the CPU 26 is arranged to normalise the pitch of the recorded utterance by determining the start and end of each pitch pulse period (illustrated in Figure 1) for example by determining the zero crossing points thereof, and then equalising the number of samples within each pitch period (for example to 140 samples in each pitch period) by interpolating between the originally stored samples.

As a result of such normalisation, the stored waveform therefore now consists of pitch pulses each of an equal number of samples. These are then stored (step 606) as the sample record 11 of the record 10 for the sound in question, to be used in subsequent synthesis.

Next, in a step 608, the linear array of samples x_0, x_1, \dots is transformed into an array of three dimensional coordinate points X_0, X_1, \dots , each coordinate point X_i corresponding to the three samples x_{i-10}, x_i, x_{i+10} , so as to embed (i.e. represent) the speech signal in a state sequence space, as illustrated in Figure 11b.

The first coordinate point is then selected (i.e. X_{10}).

The trajectory of points through the state sequence space is, as discussed above in relation to Figures 3 and 4, substantially repetitive. Thus, the trajectory consists, at any point, of a number of close "strands" or "tracks", each consisting of the equivalent portion of a different pitch pulse.

Referring to step 610, for the selected point (in this case, the first point, X_i), there will be other points on other tracks of the attractor, which are close in state sequence space to the selected point X_i . For example, as
 5 shown in Figure 11c, points X_{13} and X_{14} on a first track, and X_{21} and X_{22} on a second track, are close to the point X_i . Accordingly, in a step 610, the CPU 26 locates all the points on other tracks (i.e. in other pitch periods) which are closer than a predetermined distance D in state sequence space (D
 10 being the euclidean, or root mean square, distance for ease of calculation). To avoid a search and distance comparison of all 10,000 stored points, the CPU 26 may examine only a limited range of points in the region of $X_{(i + k.140)} \pm 5$ points, where k is an integer, so as to examine roughly corresponding
 15 areas of each pitch pulse to that in which the reference point X_i is located.

Having located a group of points on other tracks than that of the reference point X_i , the CPU 26 then stores a neighbourhood array B_i of vectors b_i , as shown in Figure 11d,
 20 in step 612. Each of the vectors b_i of the array B_i is the vector from the reference point X_i to one of the other neighbouring points on a different track of the attractor, as shown in Figures 11 and 13. A set of such vectors, represented by the neighbourhood matrix B_i , provides some
 25 representation of the local shape of the attractor surrounding the reference point X_i , which can be used to determine how the shape of the attractor changes as will be described further.

Next, in step 614, the CPU 26 selects the next point X_{i+1} along the same track as the original reference point X_i .

30 Next, in step 616, the CPU 26 progresses forward one point on each of the other tracks of the attractor, so as to locate the corresponding points on those other tracks forming the new neighbourhood to the new reference point X_{i+1} , in step 616. In step 618, the CPU 26 calculates the corresponding
 35 neighbourhood array of vectors B_{i+1} .

Because the pitch pulses of the recorded utterance differ slightly one from another, the corresponding tracks of the

stored attractor trajectory will also differ slightly one from another. At some points, the tracks will be closer together and at some points they will be more divergent.

Thus, the new set B_{i+1} of offset vectors b_{i+1} will have
 5 changed position, will have rotated somewhat (as the attractors form a loop), and will also in general be of different lengths to the previous B_i set of vectors b_i . Thus, in progressing around the attractor track from one sample to the next, the set B_i of vectors b^1_i, b^2_i (and hence the shape
 10 of the attractor itself which they represent) are successively transformed by displacement, rotation and scaling.

Next, in step 620, the transformation matrix T_i which transforms the set of vectors B_i defining the attractor in the neighbourhood of point X_i to the set of vectors B_{i+1} defining
 15 the neighbourhood of the attractor in the region of the reference point X_{i+1} is calculated in step 620. The matrix is therefore defined as:

$$B_{i+1}^T = T_i B_i^T$$

This can be rearranged to the following form:

20 $T_i^T = B_i^{-1} B_{i+1}$

In general, since B_i is a $dx3$ matrix (where d is the number of displacement vectors used, which may be greater than 3) B_i will not have an exact inverse B_i^{-1} , but the pseudo inverse can instead be calculated, as described in Moore and
 25 Penrose, "A generalised inverse for matrices", Proc. Camb. Phil. Soc., Vol. 51, pages 406-413, 1955.

The 3×3 transform matrix T_i thus calculated is an approximation to the transformation of any one of the vectors making up the neighbourhood matrix B . However, since the
 30 neighbourhood in the state sequence space is small, and since speech is locally linear over small intervals of time, the approximation is reasonable.

Next, in step 622, the CPU 26 selects the next point X_{i+1} as the new reference point and returns to step 610.

35 Thus, after the process of steps 610 to 622 has been performed for each of the points X_i corresponding to digitised speech sample values x_i , all the transform matrices thus

calculated are stored (step 624) associated with the respective data values x_i corresponding to the reference points X_i for which the matrix was derived, in a data record 12.

5 Thus, at the end of the process of Figure 12, the stored transform matrices T_i each represent what happens to a displacement vector b_i , from the point on an attractor for which the transform matrix was calculated to another point in space close by, in moving one sample forward in time along the
10 attractor. It will therefore be understood how the use in Figure 7 of the transform matrices thus calculated enables the construction of a new synthesised point on the attractor, using a stored actual trajectory forming part of the attractor, a previous synthesised point (and hence a previous
15 vector from the stored trajectory to that previous synthesised point) and the transformation matrix itself.

 The above description relates to the derivation of stored data for synthesis of a voiced sound. For storage of data relating to unvoiced sounds, only steps 602 and 606 are
20 performed, since the storage of the transform matrix is not required.

 Having derived the necessary data for each voiced or unvoiced sound in the phonetic alphabet as described above, the stored data are transferred (either by communications link
25 or a removable carrier such as a floppy disk) to the memory 8 of synthesis apparatus of Figure 6.

Reproduction of unvoiced sounds

 Mention is made in step 530 of reproduction of unvoiced sounds. As discussed above, unvoiced sounds do not exhibit
30 stable low dimensional behaviour, and hence they do not follow regular, repeating attractors in state sequence space and synthesis of an attractor as described above is therefore unstable. Accordingly, unvoiced sounds are produced in this embodiment by simply outputting, in succession, the stored
35 waveform values x_i stored for the unvoiced sound to the DAC 4. The same is true of plosive sounds.

Progression to sounds

In relation to steps 504, 524 and 528 of Figure 7, mention was made of progression to or between sounds. One possible manner of progression, usable with the above
 5 described embodiment, will now be disclosed in greater detail.

Referring to Figures 14 and 15, Figure 14 illustrates the steps making up step 504 or step 528 of Figure 7, whereas Figure 15 graphically illustrates the effect thereof.

Broadly speaking, the present invention interpolates
 10 between two waveforms, one representing each sound, in state sequence space. The state space representation is useful where one or both of the waveforms between which interpolation are performed are being synthesised (i.e. one or both are voiced waveforms). Broadly speaking, in this embodiment, the
 15 synthesised points in state space are derived, and then the interpolated point is calculated between them; in fact, as discussed below, it is only necessary to interpolate on one co-ordinate axis, so that the state space representation plays no part in the actual interpolation process.

20 The interpolation is performed over more than pitch pulse cycle (for example 10 cycles) by progressively linearly varying the euclidean distance between the two waveforms in state sequence space.

Thus, as indicated in Figure 14, the coordinates of a
 25 given point \underline{S}^c_i during transition between voiced sounds are derived from the coordinates in state sequence space of a synthesis point on the attractor of the first sound \underline{S}^a_i and a corresponding point on the attractor of the second sound \underline{S}^b_i .

30 In more detail, referring to Figure 14, in a step 702, an index i is initialised (e.g. at one).

In step 704, the current value of the synthesised attractor on the first waveform \hat{S}^a_i is calculated, as disclosed above in relation to Figure 7.

35 In a step 706, the CPU 6 scans the recorded sample values for the second sound to be progressed towards and locates (for example by determining the zero crossing points) the sample \underline{S}_i^b at the same relative position within a pitch period of the

second waveform as the point \underline{S}_i^a . In other words, if the point \underline{S}_i^a on the first waveform is the 30th point within a pitch period of the first sound from the zero crossing thereof, the point \underline{S}_i^b is also selected at the 30th point after the zero crossing of a pitch period of the second sound.

Then, a synthesised attractor point $\hat{\underline{S}}_i^b$ is calculated as disclosed above in relation to Figure 7.

Next, in a step 708, the coordinates of an interpolated point \underline{S}_i^c are calculated by linear interpolation, in step 708. It is only necessary to calculate one dimension of the interpolated attractor, since it is only the current output sample value which is desired to be synthesised, not the sample value ten samples previously or ten samples in the future. Thus, in step 708, the interpolation calculation actually performed is:

$$x_i^c = \frac{i.x_i^a + (N-i).x_i^b}{N}$$

Where N is the number of samples over which interpolation is performed.

Then, in step 709, the CPU outputs x_i^c , the current time sample thus calculated, to the DAC for and hence loudspeaker 2 for synthesis.

In step 710, the CPU 6 tests whether the end of a predetermined transition duration has been reached (e.g. of 40 samples, so that $N=40$) and, if not, in step 712 the index i is incremented, and in step 714 the next values of the synthesised attractor ($\hat{\underline{S}}_i^a$) is calculated and the next corresponding value of the attractor of the new sound $\hat{\underline{S}}_i^b$ is derived. Step 708 is then repeated.

When the last sample of the transition, $i=N$, has been reached in step 710, the CPU 6 proceeds with step 506 or step 530, as discussed above in relation to Figure 7, to synthesise the new sound corresponding to the attractor \underline{S}^b .

The above described process applies equally where a transition is occurring from silence to a stored representative sound. In this case, rather than calculating

a value for \hat{s}_i^a , the CPU 6 reads a corresponding value of zero, so that the corresponding effect is simply a linear fade to the required synthesised sound.

Likewise, when the transition is from a sound to silence, as in step 524, the same sequence as described above in relation to Figure 14 is performed except that instead of calculating successive synthesised values of the attractor of the second sound, the CPU 6 is arranged to substitute zero values, so as to perform a linear fade to silence.

10 Progression to and from unvoiced sounds

The process of progression described above in relation to Figure 14 is modified in relation to progression to or from an unvoiced sound, because rather than synthesising the unvoiced sound, the actual stored value of the unvoiced sound is reproduced. Accordingly, in progression from one unvoiced sound to another, state sequence space plays no part since it is merely necessary to interpolate between corresponding successive pairs of points in the old unvoiced sound and the new unvoiced sound. Likewise, in progressions between an unvoiced sound and silence, a linear fade to or from the value of successive points of the unvoiced sound is performed.

Second embodiment

Rather than storing the transformation matrix for each point, in the second embodiment the transformation matrix is calculated directly at each newly synthesised point; in this case, the synthesizer of Figure 6 incorporates the functionality of the apparatus of Figure 10. Such calculation reduces the required storage space by around one order of magnitude, although higher processing speed is required.

30 In this embodiment, rather than interpolating between stored sample values directly to produce output sample values as described above in the first embodiment, it is possible to interpolate to produce intermediate attractor sequences and corresponding transformation matrices describing the dynamics of the intermediate transformation sequences. This gives greater flexibility, in that it is possible to stretch the production of the intermediate sounds over as long a period

as is required.

Referring to Figure 16, in this embodiment, in a step 802, a first counter i is initialised. The counter i sets the number of intermediate templates which are produced, and is conveniently of a length corresponding to several pitch cycles (in other words, N , the maximum value for i , is around 300-400).

In a step 804, the value of another constant k is initialised; this corresponds to the number of stored points on each of the two stored waveforms (and is thus typically around 10,000).

In a step 806, a corresponding pair of points \underline{s}_k^a , \underline{s}_k^b are read from the stored waveform records 10; as described in the first embodiment, the points correspond to matching parts of the respective pitch pulse cycles of the two waveforms.

Next, in a step 808, an interpolated point \underline{s}_k^c is calculated as described in the first embodiment.

If the last point on the waveforms has not been reached (step 810), in step 812 the value of the counter along the waveforms, k , is incremented and steps 806-810 are repeated.

Thus, following execution of steps 804-812 for each stored point, around half a second of an intermediate waveform which defines a repetitive trajectory in space will have been calculated.

Then, in step 814, the CPU 6 performs the steps 610-622 of Figure 12, to calculate the transform matrices T_k for each point along this stored track.

After performance of step 814, sufficient information (in the form of a stored interpolated trajectory and stored interpolated transformation matrices) is available to synthesise a waveform of any required length from this intermediate trajectory. In fact, however, this calculated data is used to derive only a single new point in state sequence space, \hat{s}_{i+1} , by transforming the previous value of \hat{s}_i which was most recently output, in step 816.

The sample value \hat{x}_{i+1} thus calculated as part of \hat{s}_{i+1} in output in step 818, and, until the end of the transition

portion has been reached (step 820), the interpolation index i is incremented (step 822) and the CPU 6 returns to step 804 to calculate the next interpolated trajectory and set of dynamics T_k , and hence the next point to be output.

5 It will be apparent that, although in the above described embodiment, each interpolated trajectory and set of transformation vectors is used only once to calculate only a single output value, in fact fewer interpolated sets of trajectories and sets of transformation matrices could be
10 calculated, and the same trajectory used for several successive output samples.

Equally, although linear interpolation has been discussed above, it would be possible to use a non-linear interpolation (describing, for example a sigmoid function).

15 Equally, whilst the process of Figure 16 has been described for producing a progression between two sounds by interpolation, it would be possible to use the process of steps 804-818 to produce a constant intermediate sound between
20 two stored sounds, thus enabling the production of intermediate vowels or other sounds from a more limited subset of stored sounds.

Other embodiments and variations

It will be apparent from the foregoing description that many modifications or variations may be made to the above
25 described embodiment without departing from the invention.

Firstly, although the foregoing describes storage of multiple pitch pulse sequences, it would be possible to store only a single pitch pulse sequence (i.e. a single track of the attractor) for each voiced sound, since the synthesis process
30 will enable the reproduction of multiple different synthesised pitch pulse sequences therefrom. This may reduce the volume of data necessary for storage under some circumstances.

Indeed, rather than storing an actual attractor track, it will be clear that some other reference curve (for example
35 an averaged attractor track) could be stored, provided that the transformation matrices from such other curve to the actual attractor strands had previously been calculated as

described above.

Although in the above described embodiment, the dynamics of the speech waveform (in the state sequence space) are described by a neighbourhood matrix describing the transformation of vectors running between adjacent strands of an attractor, it will be clear that the transformation matrix could instead describe the evolution of a point on the attractor directly.

However, we have found that describing the transformation of a difference vector between an actual attractor and another actual or synthesised attractor has the virtue of greater stability, since the synthesised waveform will always be kept reasonable close to an actual stored attractor.

Rather than progressing between respective synthesised values of voiced sounds, it is possible to progress between respective stored values, in the same manner as described above in relation to progress between unvoiced sounds; in this case, the progression is therefore simply performed by linear interpolation between successive pairs of corresponding stored sample points of the two sounds, although an improvement in performance is obtained if the interpolation is between points from corresponding portions of pitch pulses as described above.

To determine corresponding points in successive pitch pulses, rather than utilising zero crossings as discussed above, it would be possible to record the physical motion of the human vocal system using a laryngograph monitoring the human speaker recording the utterances as described in relation to Figure 12, to directly identify corresponding physical positions of the human vocal system. Equally, the positions in state sequence space of the respective attractors of the two sounds could be used to identify respective portions of the sounds (although this method can lead to ambiguities).

The speech synthesizer of the embodiment of Figure 6 is described as generating samples one by one at the time each sample is calculated, but it would of course be possible to

generate and buffer a sequence of samples prior to reproduction.

It would be straightforward to modify the synthesizer disclosed above in relation to Figure 6 to provide that the
5 CPU effects amplitude control by scaling the value of each output sample calculated, or by direct control of an analog amplifier connected to the loudspeaker 2.

In this case progressions to and from silence may additionally or alternatively utilise a progressive amplitude
10 increase or reduction.

Equally, it would be straightforward to provide for variation of pitch in the described embodiment, by altering the rate at which the CPU 6 supplies output samples to the digital to analog converter 4.

15 Although in the above described embodiment a digital to analog converter and a loudspeaker are provided, it is of course possible for the digital to analog converter and loudspeaker to be located remotely. For example, the speech synthesizer may in another embodiment be provided at a site
20 within a telecommunications network (for example at a network control station or within an exchange). In such a case, although the speech synthesizer could provide an analog output, it may equally be convenient for the speech synthesizer to supply a train of digital sample outputs since
25 the speech carried by the telephone network may be in digital form; eventual reconstruction to an analog waveform is therefore performed in this embodiment by local exchange or end user terminal components rather than a digital to analog converter and loudspeaker forming part of the speech
30 synthesizer. For example, such an embodiment may be applied in relation to automated directory enquiries, in which stored subscriber telephone number digital information is reproduced as a speech signal under the control of a human operator or a speech recogniser device.

35 It will be apparent that many other modification and variants may be formed without departing from the essence of the present invention.

CLAIMS:

1. A method of generating a synthetic waveform output corresponding to a sequence of substantially similar cycles, comprising the steps of
 - 5 (a) generating a synthetic waveform sample;
 - (b) generating a successive waveform sample from said synthetic waveform sample and data defining the transformation followed by said cycles in the temporal vicinity of said synthetic waveform sample;
 - 10 (c) designating said successive waveform sample as a synthetic waveform sample and repeating step (b);
 - (d) repeating step (c) a plurality of times to generate a sequence of said successive waveform samples corresponding to a plurality of said cycles; and
 - 15 (e) outputting the samples of said sequence to generate a waveform.
2. A method according to claim 1, in which said waveform comprises voiced speech.
3. A method according to claim 1 or claim 2, in which
 - 20 said data defining said transformation does so by reference to a predetermined reference waveform sequence.
4. A method according to claim 3, in which said reference waveform sequence comprises a stored speech waveform.
- 25 5. A method according to any preceding claim, in which said steps (a) and (b) comprise generating a plurality of values representing said waveform sample values as a point in a multidimensional space in which corresponding portions of successive said cycles are substantially superposed.
- 30 6. A method according to claim 5 when appended to claim 3 or claim 4, in which the transformation approximates that

which would transform a first displacement vector, extending from a first time point on said reference waveform sequence to a corresponding time point on the waveform to be synthesised, to a second displacement vector extending from
5 a second point, successive to the first, on said reference waveform sequence to a corresponding second point on the waveform to be synthesised.

7. A method according to any of claims 3 to 6, in which a given successive waveform sample is derived in accordance
10 with data from a point on said reference waveform sequence at a position within a said cycle which corresponds to that of said given successive waveform sample, and at least one other point on said reference waveform sequence offset in time therefrom.

15 8. A method according to any preceding claim, in which said step (b) comprises calculating said transformation from a set of stored waveform values.

9. A method according to any preceding claim in which the initial performance of said step (a) to initial synthesis
20 of said waveform comprises a step of selection of an initial value which differs from a previous initial value selected on a previous synthesis of said waveform.

10. A method according to claim 9 in which said selection step comprises applying a pseudo random number
25 generation algorithm to select said value.

11. A method according to claim 9 or claim 10 in which said step of selection comprises referring to a stored waveform sample value and calculating a synthesised initial waveform value similar but different to said stored waveform
30 value.

12. A method of synthesis of a voiced speech sound

comprising calculating each new output value from the previous output value using data modelling the evolution, over a short time interval, of the voiced speech sound to be synthesised.

13. A method of concatenating two cyclical sounds,
5 comprising progressively interpolating between pairs of values of said sounds at corresponding points within the cycle of each of said sounds.

14. A method of synthesising a cyclical sound
intermediate between two other cyclical sounds, for each of
10 which a succession of sample values corresponding to a plurality of cycles are stored, comprising the steps of generating interpolated waveform samples consisting of a succession of values each of which is interpolated from a pair of points, one each respectively from corresponding portions
15 from a cycle of each of the stored waveforms; generating a model of the evolution, over a short time interval, of the interpolated waveform; and calculating each successive output value from a previous output value using said evolution model.

15. Synthesis apparatus arranged to perform the method
20 of any preceding claim.

ABSTRACT

A method of synthesis of a voiced speech sound comprising calculating each new output value from the previous output value using data modelling the evolution, over a short time
5 interval, of the voiced speech sound to be synthesised.

(Figure 7)

1/12

FIG. 1







FIG. 2



2/12



3112


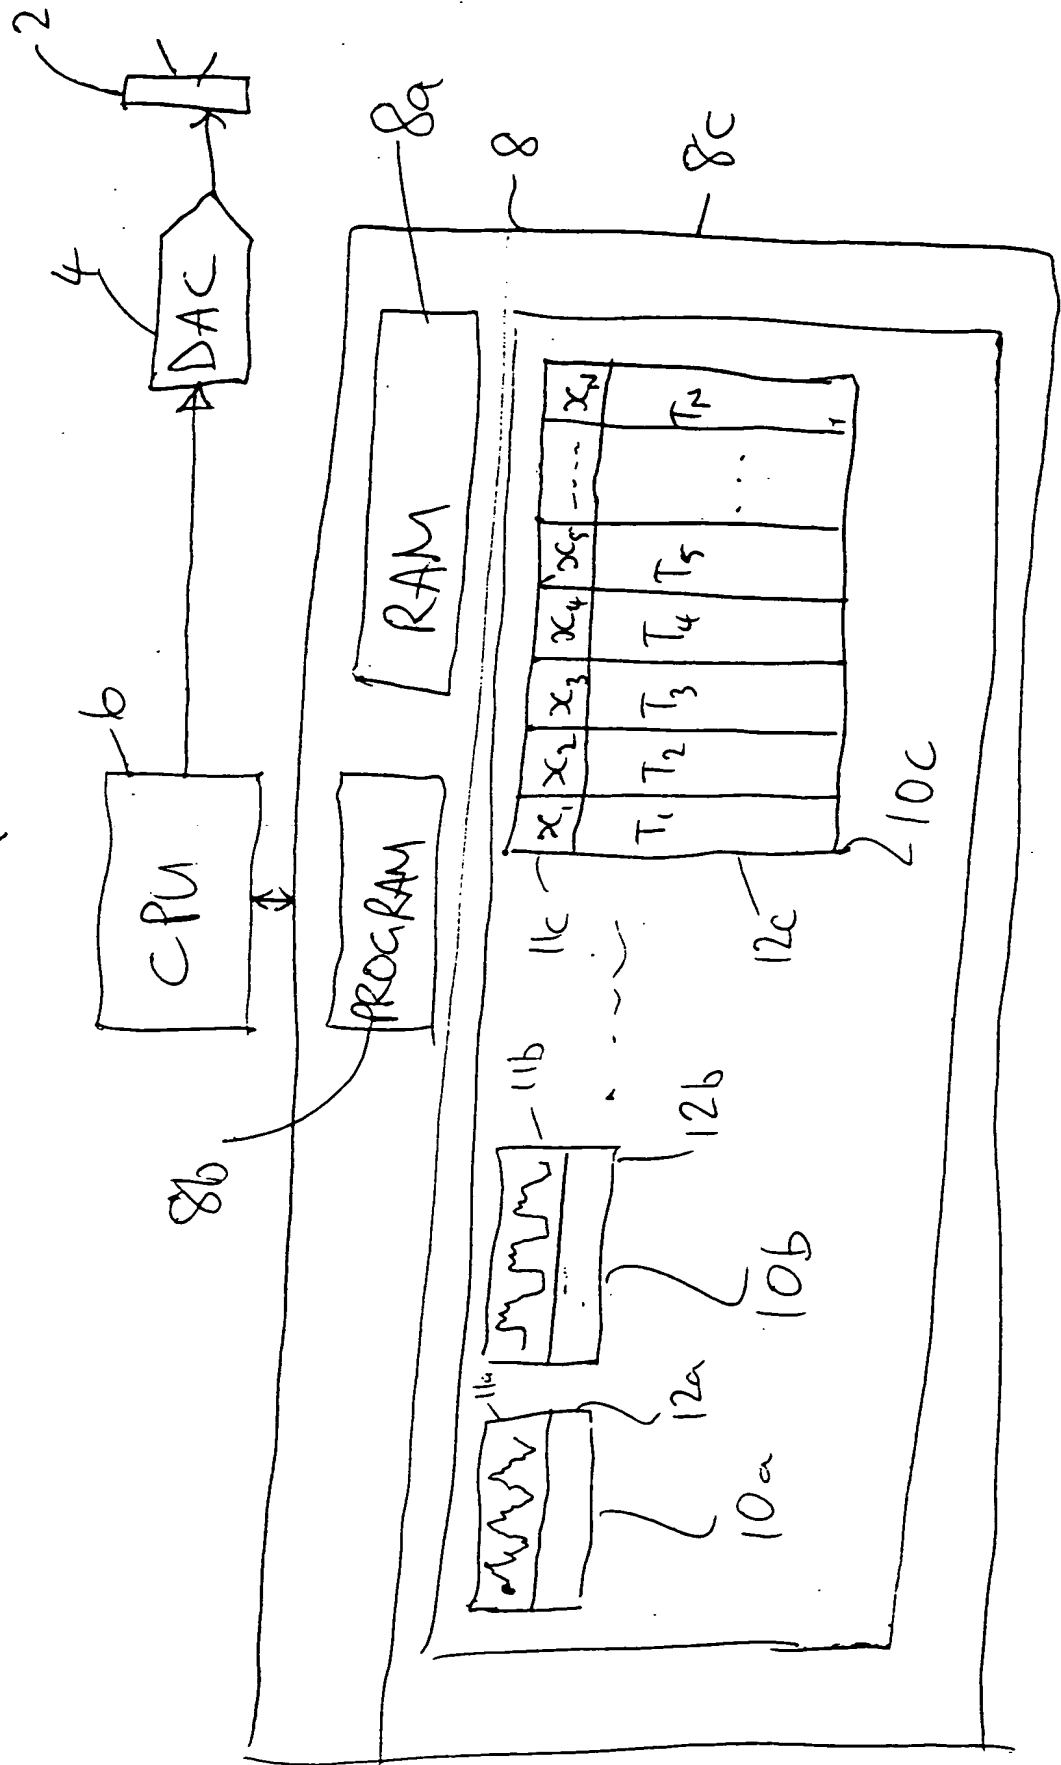


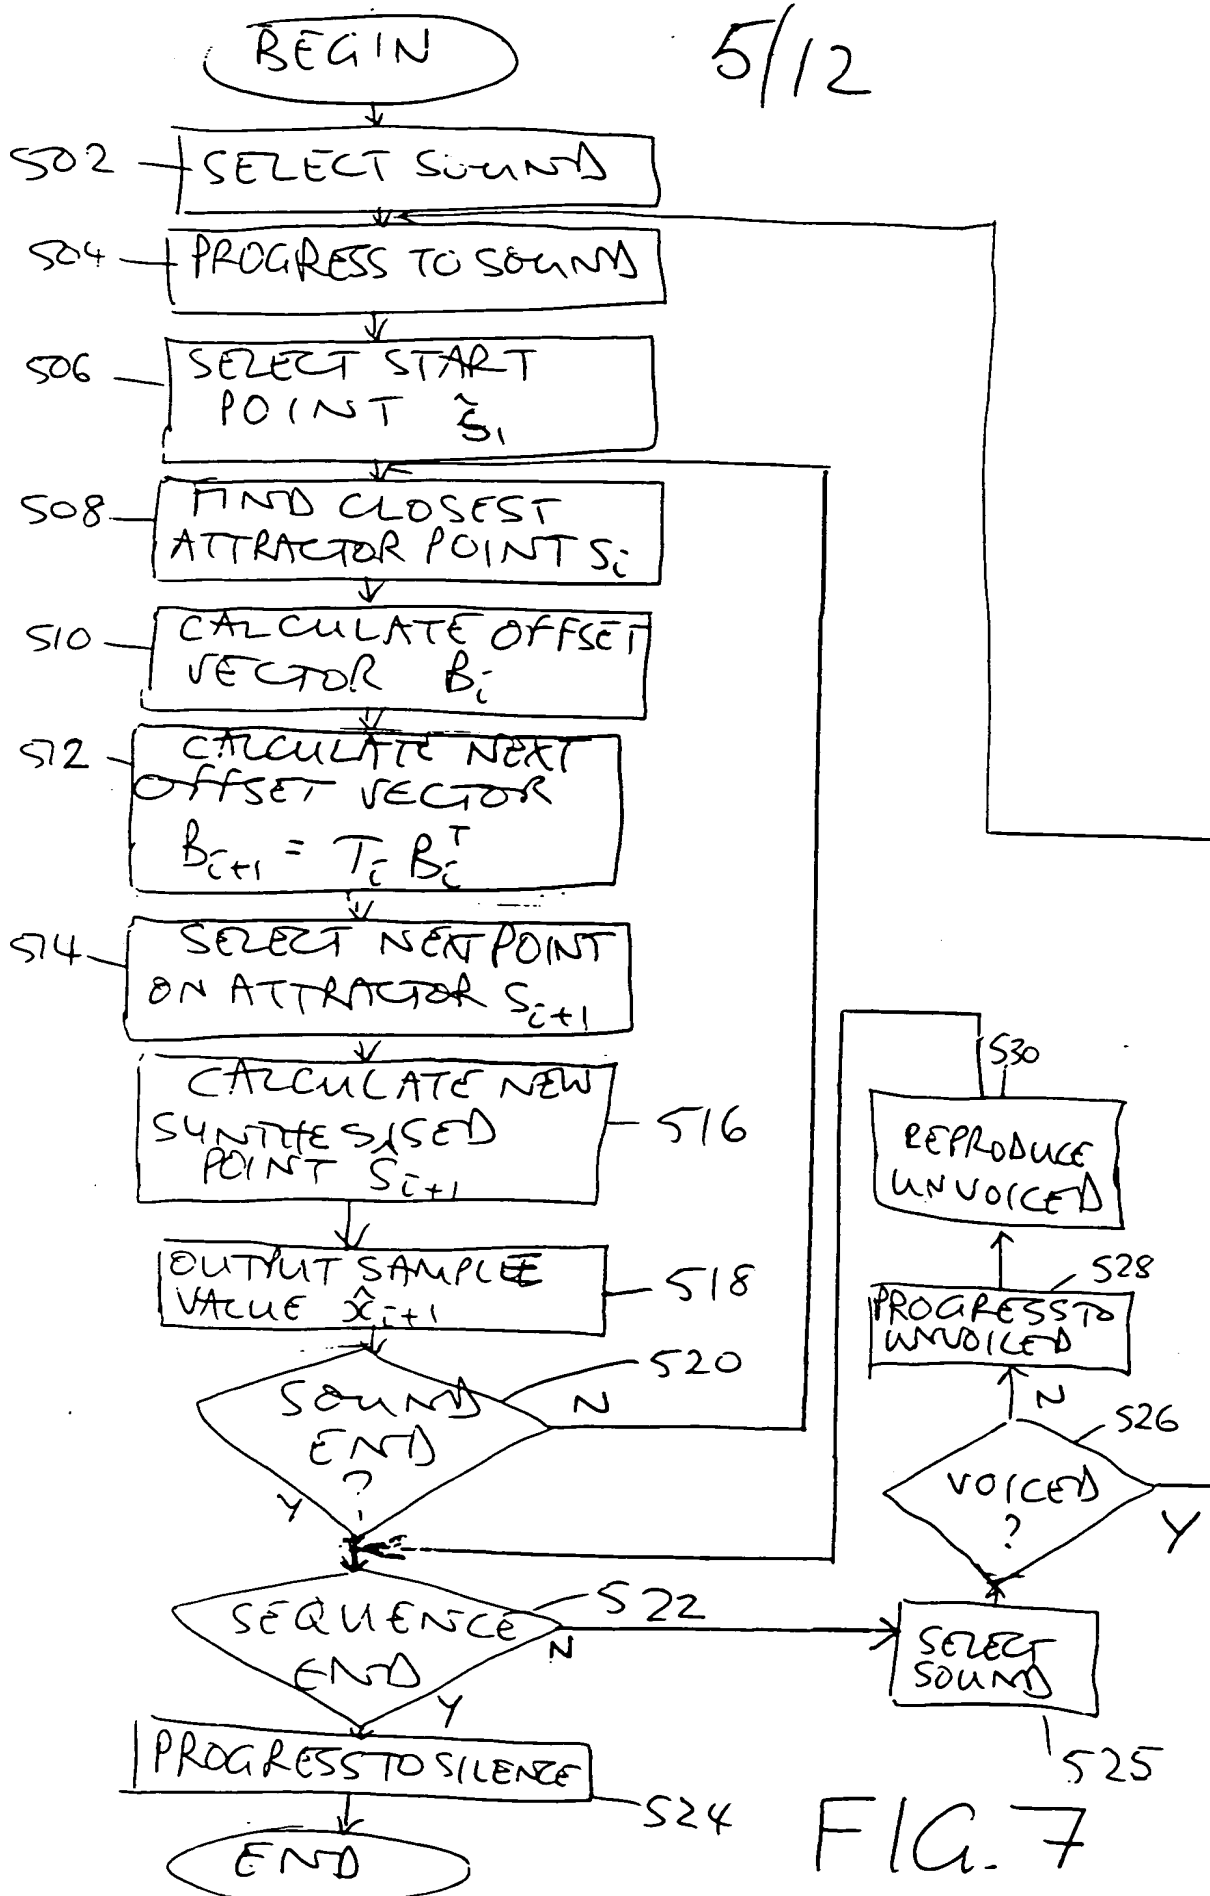
FIG. 5 f_1

4/12

FIG. 6



5/12



6/12




FIG. 8




FIG. 9c

- x Synthesis start point \hat{S}_i
- Nearest point S_i

FIG. 9d

7/12

FIG. 9a





FIG. 9b






FIG. 11




FIG. 12

9/12





FIG. 10






FIG. 1A

11/12




FIG. 15